

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication: 06.03.1996 Bulletin 1996/10	(51) Int Cl. ⁶ : G06F 12/02
(21) Application number: 95480098.3	
(22) Date of filing: 24.07.1995	
(84) Designated Contracting States: DE FR GB	(72) Inventors: <ul style="list-style-type: none">• Munroe, Steven Jay Rochester, Minnesota 55901 (US)• Sirjani, Abolfalz Austin, Texas 78729 (US)• Voldal, Erik Edward Rochester, Minnesota 55901 (US)
(30) Priority: 30.08.1994 US 299042	
(71) Applicant: International Business Machines Corporation Armonk, N.Y. 10504 (US)	(74) Representative: Lattard, Nicole Compagnie IBM France Département de Propriété Intellectuelle F-06610 La Gaude (FR)

(54)

System and method combining a global object identifier with a local object address in a single object pointer

(57) An object pointer data structure for efficiently combining an object identifier and an object address for use in object-oriented programming systems. An object address is a value that allows a client application or process to conduct high-performance operations on an object in the client's local virtual address space. An object identifier is a value that can be used to uniquely identify an object for the lifetime of that object across some defined domain, such as an entire universe of computer systems. The data structure of this invention defines an object pointer that is larger than the object address but smaller than the combination of the object identifier and object address. The truncated object pointer structure preserves all information from both object address and object identifier by forcing a portion of the local object address in each address space to be equal to a portion of the invariant object identifier. A local pointer mapping table may be used for efficiency in assigning local addresses to restored objects in each process.




FIG. 8

Description

BACKGROUND OF THE INVENTION

1. Field of the Invention

This invention relates generally to addressing objects in an object-oriented programming environment and more particularly to a system and method for combining a global object identifier with a local object address in a single object pointer.

2. Discussion of the Related Art

Large computer systems permitting many users to simultaneously use single computer resources are herein denominated multitasking computer systems. Such multitasking systems can be interconnected through networks to permit concurrent processing over a distributed plurality of computer resources. The concurrent activities of a large number of simultaneous users are managed by an operating system, which must provide the necessary data structures for tracking and controlling many concurrent processes. Requisite data structures include those needed to manage the distributed memory and the multiple central processing unit (CPU) resources that are used by user processes throughout the network.

If every user process were completely independent, having its own dedicated resources, and there were no concerns about which resources each process must use, operated systems could be relatively simple. However, networked or distributed computer resources are usually shared by many user processes, each requiring access to commonly-owned resources. In fact, each user process may generate a number of simultaneous execution threads that must also share resources among themselves and communicate with the threads and processes of other users.

Modern operating systems employ object architectures to govern data access and transfers among concurrent processes. As used herein, objects include data structures holding information that must be protected from unauthorized access by other user processes. Objects are created by the operating system responsive to user requests and are structured to permit access by the requesting user through system routines that protect the integrity of each object.

The concepts of object identifier and object address (or reference) are well-known in the typical object-oriented programming art. As used herein, an object identifier denominates a value that can be used to uniquely identify an object for the lifetime of that object across some defined domain, such as a single system, a network of systems or the set of all systems in some predefined universe. As used herein, an object address (or reference) denominates a value that allows a user application to conduct high-performance operations on an object in a local virtual address space assigned to the associated

user application process. An object address need not be unique in any domain broader than the local virtual address space in which it is used. As known in the art, the concepts of object identifier and object address are separate and distinct.

An object identifier is needed to identify an object having a lifetime that exceeds its local address lifetime. Object identifiers are used, for instance, to manage a network of objects, herein denominating a structure embracing a multiplicity of objects linked together by stored object identifiers each linked to an object within the multiplicity. When an object must be moved from one system to another or from one process to another, for instance, to permit controlled access by many different users, the only consistent means available for tracking such objects is the object identifier, which remains invariant during the interprocess transition from one virtual address space to another.

Because object address and object identifier are conceptually distinct and have independent values, the systems programmer must provide for both. For instance, storage must be allocated for both, storing and updating operations must be allocated for both and so forth. This doubled complexity reduces programming productivity and increases the probability of programming error. There is, accordingly, a clearly-felt need in the art to provide both functions in a single data structure. Although this could be accommodated merely by concatenating the two values into a single double-sized data structure, such a solution offers no reduction in programming complexity because no storage allocation or updating operations are thereby avoided.

Practitioners in the art have proposed a number of solutions to the object-oriented programming complexity problem. For instance, one solution uses a server process that temporarily impersonates the characteristics of a client process when the client process performs a remote procedure call on the server process in a multitasking system. This impersonation is accomplished by creating a new object identifier list that is either the same as the client process list or represents the union of the server and client lists. Such an access control list is provided in each object to define those identifiers required to access the object. Thus, when the identifier list is modified to map all necessary identifiers to local addresses, access checking software in the operating system may then enable the impersonating process to access the specified object. This solution arguably increases programming complexity to accomplish the necessary manipulation of mapping tables that relate object identifiers to local addresses.

Similarly, another solution uses a system for creating "conditional" objects having different object pointers for accessing selected system resources. This object-based operating system supports multiple levels of visibility, thereby permitting objects to be manipulated only by processes within the object's "range" of visibility. An object or an object network can be moved up to a

higher visibility level for more object sharing and down to a lower visibility level for more restricted or privileged access. The visibility restriction is achieved by giving each process a pointer (object identifier) to its own process-level container directory and providing for a "public display" container directory permitting interprocess object access. Thus, such a scheme requires a multiplicity of different directories and object identifier mapping tables within each process for linking each global object identifier to various local virtual addresses, increasing rather than reducing programming complexity.

Another approach achieves a virtual look-aside facility for maintaining named data objects in class-related data spaces in virtual storage by providing for an ordered list of "major names" within each user process that is sequentially searched for a specified "minor name" to obtain a virtual storage copy of a data object. This requires a control structure designed to capture information about the location of data within a tree structure that is implicitly created by a naming technique that is commonly used during routine caching operations. This information by-product, selectively used to improve performance in subsequent transactions, does nothing to reduce the programming complexity associated with object naming conventions.

Also known in the art is an object identifier generator for producing unique identifiers in a distributed computer system by concatenating the identifier of the creating node firstly with a creation time from an associated clock adjusted to ensure uniqueness, secondly with a random name-sequence, and finally with a system version number. This generator ensures the creation of an object identifier that is unique across an indefinite universe of computer systems but does not affect the programming complexity problem.

Other related approaches include a system environment that employs workspace name tables for linking packaged workspaces (user processes). An active workspace accesses a loaded copy of the packaged workspaces using the external names of named objects in a packaged workspace rather than the internal names of its own named objects. This system is controlled through a packaged workspace name table stored in the active workspace, which can be merged with other name tables from other packaged workspaces and also shared among a plurality of active workspaces. As with other known object identifier to local address mapping schemes, this system merely maps object identifiers to local addresses and does nothing to reduce programming complexity.

It can be appreciated from the above survey that practitioners in the art, until now, appear to rely on mapping tables with controlled access and duplication privileges to establish and control the relationship between global object identifiers and local virtual addresses, which does nothing to improve programming efficiency. The related unresolved deficiencies are clearly felt in the art and are solved by this invention in the manner de-

scribed below.

SUMMARY OF THE INVENTION

The system of this invention introduces a new data structure and a process that preserves part of the object identifier in each new local address selected for an object that is transferred into a new virtual address space. By preserving this part of the object identifier, a single truncated pointer is then sufficient to carry all object identifier and object address information. A local mapping table may be used to manage local address assignments but is not essential to the system of this invention.

It is an object of the system of this invention to provide an efficient mechanism for combining an object identifier and an object address into a single object pointer that provides the features of both. It is an advantage of the system of this invention that a single object pointer is sufficient for both global object identifier and local object address with less complexity and in less space than is required separately for both.

It is another object of the system of this invention to provide such an object pointer that can be controlled and manipulated by the operating system as a single data structure to reduce programming complexity. It is a feature of the data structure of this invention that both object identifier and object address information is contained in one object pointer that appears to the operating system to be a single data structure.

The foregoing, together with other objects, features and advantages of this invention, can be better appreciated with reference to the following specification, claims and the accompanying drawing.

BRIEF DESCRIPTION OF THE DRAWING

For a more complete understanding of this invention, reference is now made to the following detailed description of the embodiments, as illustrated in the accompanying drawing, wherein:

Fig. 1 is a functional block diagram showing a multitasking computer system from the prior art;

Fig. 2 is a functional block diagram illustrating the virtual memory spaces of several exemplary concurrent processes;

Figs. 3A-3B illustrate the object identifier and object address data structure parsing of this invention;

Figs. 4A-4B show an exemplary embodiment of the hybrid pointer data structure of this invention;

Fig. 5 illustrates an exemplary object mapping table embodiment of this invention;

Fig. 6 is a functional block diagram illustrating the

object pointer mapping control procedure of this invention;

Fig. 7 is a functional block diagram illustrating the lazy object pointer mapping procedure of this invention for networked objects; and

Fig. 8 is a functional block diagram of an exemplary embodiment of the object pointer system of this invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

The system of this invention is an object-oriented programming structure suitable for application to multitasking systems. In general, object-oriented programming structures include data structures and procedures that control user-definable objects. As used herein, "objects" are data structures holding information that is used by the operating system or by user applications. For instance, a "process object" is a particular object type known in the art for storing information required by the operating system to track the status of a particular user process.

A "process" is the entity to which a virtual memory address space is assigned by the operating system and also is the entity to which process-level objects are assigned. As is known in the art, a particular user may employ multiple processes simultaneously. Whenever a particular user demands system resources, a "top level" process is created to manage all related activities. Any process, including any user process or the top level process, may cause the creation of additional processes, denominated subprocesses or "child" processes. Thus, any "parent" process may create one or more "child" processes.

Fig. 1 illustrates a computer system 10 that is typical of multitasking computer networks known in the art. System 10 includes a high-speed central processing unit (CPU) 12 that concurrently runs the several processes 14, 16, 18 and 20. CPU 12 may be either a single processor or may include multiple distributed processors. Each process 14-20 is associated with its own virtual memory space, which is mapped in part into high-speed primary memory 22 and in part into lower-speed secondary memory 24 by a virtual memory manager 26. Each process 14-20 is allocated a portion of the available computer resources, including selected peripheral devices such as terminals 28 and 30 and other input-output (I/O) devices 32 and 34. Process allocations also include specified sets of data and data structures in the distributed memory 22 and 24.

Operation and resource allocation in computer system 10 is supervised by the operating system 36 which resides in primary memory 22 for most purposes. Multitasking computer system 10 may also include one or more remotely-located computer systems, exemplified

by remote computer 38, which may include other terminals, I/O devices, primary and secondary memory devices and the like (not shown). Data communications among CPU 12, peripheral devices 28-34 and remote computer 38 are controlled by a bus interface 40 in one of several fashions known in the art.

Fig. 2 shows a functional block diagram illustrating the allocation of virtual memory space for several concurrent user processes 14-18. The virtual memory space associated with each process includes a virtual address space, exemplified by virtual address space 42 for process 14, that can be accessed by "user mode" programs as well as "kernel mode" programs having system-wide privilege, including "executive" mode 44, "kernel" mode 46 and "hardware" mode 48 privilege levels known in the art. The virtual memory space 44-48 allocated to kernel mode is common to all user processes running in the multitasking computer system 10. Thus, a predefined portion (44, 46 and 48) of the virtual memory space associated with each user process is occupied by operating system 36 and its data structures. The user mode portion 42 occupies the remainder of the associated virtual memory space assigned to the user process 14.

When a user mode program, exemplified by a user call 50 in user process 14, creates an object or performs an operation on an object, user process 14 calls a kernel mode routine 52 in the executive portion 44 of the associated virtual memory space to perform the necessary object-creation operations. After completion of kernel routine 52, control is returned to user program 50 in process 14. Thus, kernel mode programs are responsible for objection creation and object transfers between independent user processes, such as an object transfer between user process 14 and 16, for instance. Thus, as is well-known in the object-oriented programming art, operating system 36 (through kernel mode programs) is privileged to move objects among user processes 14-20 even though each such user process is associated with a completely independent virtual address space, exemplified by virtual address space 42.

Fig. 3A illustrates an object identifier 54 that is created by concatenating an address space identifier 56, representing the virtual address space in which the identified object was first created, to a local object address 58, representing the first virtual address for the identified object. This particular known method for creating object identifier 54 is arbitrary but useful because it ensures creation of a unique object identifier for each new object in system 10.

Fig. 3B illustrates a local object address 60 that is created by concatenating a segment identifier 62 in a virtual address space with an offset 64, representing the precise storage location of the object within segment 62 of the associated virtual address space.

Fig. 4A illustrates the data structure of the object pointer 66 of this invention. By defining pointer 66 to be larger than object address 60, enough information can be stored in pointer 66 to identify both object address 60

and object identifier 54. Instead of using the simple method of appending object identifier 54 to object address 60, which forces object pointer 66 to become unnecessarily large, the unexpectedly advantageous solution of this invention involves preserving some of object identifier 54 in object address 60. This solution opposes the usual practice of preserving address 60 in identifier 54, and is illustrated in more detail in Fig. 4B, which summarizes the various parsings of this invention for structures 54, 60 and 66.

Although Fig. 3B shows object address 60 to be a virtual address logically divided into segment ID 62 and offset 64, and Fig. 3A shows object identifier 54 to be logically divided into address space identifier 56 and local object identifier 58, the method of this invention imposes a more general and useful parsing on these particular structures. This generality can be appreciated from Fig. 3A wherein object identifier 54 is divided into a first part (I) 68 and a second part (J) 70. Similarly, in Fig. 3B, object address 60 is divided into a first part (A) 72 and a second part (B) 74. Parts (I) 68 and (J) 70 do not necessarily align with identifiers 56 and 58. Similarly, parts (A) 72 and (B) 74 do not necessarily align with segment identifier 62 and offset 64.

Pointer 66 includes three or four parts. These are a first part (P) 75 corresponding to first part (I) 68 from object identifier 54; a second part (Q) 76 corresponding to first part (A) 72 from object address 60; a third part (R) 77 corresponding both to second part (J) 70 from object identifier 54 and second part (B) 74 from object address 60; and an optional fourth part (S) 78, which is useful for internal object pointer flagging purposes as discussed below.

Considering, as a specific example, the IBM AS/400 addressing architecture known in the art, segment ID 62 is 40 bits long, offset 64 is 24 bits long, and address space identifier 56 can be a 32-bit value that is sufficient to uniquely identify the virtual address generator that produced the original local object address 58, which is $40+24=64$ bits in length. Thus, object identifier 54 is $64+32=96$ bits or 12 bytes (12B) in length. Object address 60 is 64 bits = 8B in length. These exemplary values imply that $I+J=12B$ and $A+B=8B$, which are generous considering that address space identifier 56 could be associated with an entire system (not just a process). A 32-bit address space identifier 56 permits the combination of up to four billion individual systems divided by the typical number of times each system is reloaded from scratch during its lifetime. Accordingly, a 12B object identifier 54 should be unique across all existing systems over all time. A more general object identifier that is unique across all heterogeneous systems, such as system 10 discussed above in connection with Figs. 1 and 2, merely requires appending an identifier field that uniquely distinguishes them from other system types.

The one byte type field (S) 78 in Figs. 4A-4B is used in the AS/400 system to distinguish various pointer types and to flag whether the pointer is "resolved" or not. In an

illustrative embodiment of object pointer 66, the object identifier prefix 79 is set to 7B in length, which sets the size of first part (I) 68. With an 8B length for object address 60, the entire length of object pointer 66 (SPQR) is $1B+7B+8B=16B$. Because object identifier prefix 79 is 7B long, the object identifier suffix 80 is $12B-7B=5B$ in length, leaving $8B-5B=3B$ for the segment identifier (SID) block 82, which sets the size of first part (A) 72 in object address 60.

The 5B overlap of object ID suffix 80 with virtual address 60 is simple to implement when an object is first created because local object identifier 58 is then identical to the object's original virtual address 60 in the base virtual address space. An important advantage of the system of this invention is the ease of managing the 5-byte overlap when an object is moved from one process to another. This is accomplished by choosing an available 8B address (A_1B_1) in the target virtual address space such that the five object identifier suffix bytes are the same as they were in the previous process. In other words, referring to Fig. 4B, object address (AB) 72-74 is chosen such that field (B) 74 is equal to field (J) 70, both of which are preserved in pointer field (R) 77. With a 40-bit long object ID suffix 80, 24 bits of freedom remain in field (A) 72 for assigning an address based on SID block 82. The 24 bits of freedom in assigning new virtual addresses can be traded off against the number of available address spaces (the size of address space identifier 56). The specific field lengths mentioned above for this embodiment are exemplary only and do not serve to limit the application of the system of this invention. For instance, object identifier suffix 80 may be smaller than SID block 82 for improved address generation flexibility.

Fig. 5 shows two processes 14 and 16 each including a mapping table exemplified by mapping table 84 in process 14. Mapping table 84 may be used to more efficiently allocate local object addresses (AB) in a particular virtual address space, as discussed below. Although many different procedures can be used to efficiently manage the selection of values for SID block 82 (field (Q) 76 of pointer 66), one useful approach is now described in connection with Fig. 6.

In Fig. 6, an object is first created at step 86 in a base process. In step 88, an address generator assigns an address (A_0B_0) to the locally-created object (IJ). For management purposes, the address generator chooses addresses starting with low-order addresses for locally-created objects, tracking which SID block values ($Q=A$) are in use for such local objects by adding an entry to mapping table 84 (Fig. 5), for instance. An object pointer ($P_0Q_0R_0$) is then created in the base space at step 90.

A requirement to move object (IJ) to a first process different from the base process initiates a system response at step 91. After copying object (IJ) to the first process in step 91, the address generator assigns a new address (A_1B_1) in the first virtual address space at step 92. This is accomplished with reference to object pointer ($P_0Q_0R_0$) by forcing $B_1=R_0$ in the first address space.

The address generator chooses A_1 in step 92 by using a rule that allocates A_1 values for objects restored from other systems, such as selecting from a list of available high-order addresses. Thus, A_1 in the first address space is selected from the top of the available address list, although the value A_0 (step 88) was selected from the bottom of the available address list in the base space. By maintaining mapping table 84 (Fig. 5) to link the object identifier prefix (I) to SID block value (A) for all objects ever present in virtual address space 14, the local address generator (discussed below in connection with Fig. 8) can quickly select the appropriate value for A_1 responsive to the object identifier prefix (I) information contained in field (P) of object pointer (PQR).

In other words, when an object is being restored, the address generator first consults table 84 to find object prefix value (I_i). If I_i is already in table 84, then the corresponding SID block value (A_i) is immediately selected from the table to complete the new storage address (AB). If object identifier prefix (I_i) is not present in mapping table 84, then SID block value (A_i) is selected from the top of an available address list and the two value entry (I_i, A_i) is added to table 84. Finally, in Fig. 6, a new object pointer ($P_1Q_1R_1$) is created and stored by forcing $P_1=P_0$ and $Q_1=A_1$ (R_1 is always equal to $R_0=B$).

When moving an object that is part of a network of objects (containing pointers to other objects) from one address space to another, the internal object pointers must also be updated to reflect assignment to a new virtual address space. Fig. 7 illustrates an exemplary procedure for handling these internal object pointer updates. Assuming that the internal pointer locations in each object being moved are known (the AS/400 system accomplishes this by appending pointer tags to pointer fields within an object), the internal pointers can easily be updated synchronously whenever restoring an object to a new address space. This is simply done by looking up object identifier prefix (I) for each internal pointer in mapping table 84 (Fig. 5) to identify a corresponding SID block value (A) and then updating the internal pointer field ($Q=A$) corresponding thereto. If object identifier prefix (I) does not appear in mapping table 84, then a new SID block value (A) is assigned and the two value entry (I, A) is added to table 84.

However, as shown in Fig. 7, internal pointers may also be updated asynchronously (lazy update) by using the fourth object pointer field (S) 76. In Fig. 7, step 86 is followed by step 96, wherein a multiplicity of internal object pointers (SPQR) are created for object in the base virtual address space. In step 98, the original object is linked to this multiplicity of other objects through the storage of internal object pointers (SPQR). Similarly, after step 91, the flag field (S) is updated for all internal pointers at step 100 by changing the flag field value to a predetermined value S_w representing restoration of the base object (IJ) to a first process. Following completion of steps 92 and 94, each internal object pointer (S_wPQR) is then asynchronously updated at step 102 to reflect

completion of the transfer to the first virtual address space. For instance, internal pointer update step 102 may be deferred until the first reference to the copied object. Step 102 may then be accomplished in the manner discussed above in connection with mapping table 84.

As an alternative to using type field (S) 76, the operating system may reserve a specific value of part (A) 72 that represents an "unresolved address". This establishes a range of virtual addressing that is never allocated by the address generator, raising a "segment fault" when referenced with that prefix. Such special value replaces part (Q) 76 when the object address is unresolved. The effect is to conserve object identifier 54 while invalidating local object address 60. When a user attempts to use the invalid local object address, the system detects a segment fault, which is then used to initiate the "lazy update" of unresolved addresses by replacing part (Q) 76 with the correct part (A) 72 value from mapping table 84. This alternative internal pointer flagging scheme can be advantageously used in any system employing a standard storage protection mechanism.

Fig. 8 provides an exemplary illustration of the data and process objects required in each virtual address space, exemplified by virtual address space 14, for the system of this invention. Mapping table 84 is discussed above in connection with Fig. 5. The local address generator 104 is also discussed above in connection with Figs. 6 and 7. An object identifier generator 106 is coupled to local address generator 104 and is responsible for assigning permanent object identifiers to objects created in virtual address space 14. Local address generator 104 is coupled to mapping table 84 and also to an available address list 108, using information from both to select addresses for new and restored objects in the manner discussed above. The object pointer generator 110 is coupled to local address generator 104 and from there to object identifier generator 106, thereby permitting the creation of object pointers as discussed above in connection with Figs. 4A-4B. Finally, an internal pointer updater 112 functions to update all pointers internal to restored objects in the manner discussed above in connection with Fig. 7. The program object and data object elements shown in Fig. 8 are also coupled to operating system 36 discussed above in connection with Figs. 1 and 2.

Claims

1. In a multitasking network of one or more computer systems with distributed memory means for storing data and data structures, including a plurality of objects each created by some one of a plurality of concurrent processes each associated with at least one of a plurality of virtual address spaces in said distributed memory means, said each object being stored upon creation in an associated said virtual

address space at a base object address (A_0B_0) with a first part (A_0) and a second part (B_0), said each object including an object identifier (IJ) with a first part (I) and a second part ($J=B_0$), a method for moving said each object to a second said process associated with a second said virtual address space, said method comprising the steps of:

(a) combining said object identifier (IJ) and said base object address (A_0B_0) to produce a base object pointer ($P_0Q_0R_0$) with a first part ($P_0=I$), a second part ($Q_0=A_0$) and a third part ($R_0=J=B_0$);

(b) associating said base object pointer ($P_0Q_0R_0$) with said each object in said distributed memory means;

(c) selecting a first object address (A_1B_1) with a first part (A_1) and a second part (B_1) in said second virtual address space such that $B_1=B_0=J$;

(d) if not already present, copying said each object to said second object address (A_1B_1) in said second virtual address space in said distributed memory means; and

(e) associating a second object pointer ($P_1Q_1R_1$) with said copied object in said distributed memory means such that $P_1=P_0=I$, $Q_1=A_1$ and $R_1=R_0=J$.

2. The method of claim 1 wherein said each process further includes, in an associated said virtual address space, an object mapping table for storing an entry (P_iQ_i) representing the first part (P_i) and the second part (Q_i) for each said object pointer created by said process and wherein said selecting step (c) comprises the steps of:

(c.1) if there exists an entry (P_iQ_i) in said object mapping table such that $P_i=P_0$, then selecting said second object address (A_1B_1) such that $A_1=Q_i$; otherwise

(c.2) adding a new entry (P_1Q_1) to said mapping table such that $P_1=P_0$ and $Q_1=A_1$.

3. The method of claim 1 or 2 wherein said distributed memory means further includes a plurality of networked objects each including one or more internal object pointers (SPQR) referring to other said networked objects, each said internal object pointer (SPQR) including a fourth part (S), said method further comprising the steps of:

(d.1) replacing said fourth part (S) of each said internal object pointer (SPQR) with a first pre-

determined value denoting pending relocation of said each object to a new process; and

(f) for each said internal object pointer (SPQR), responsive to said fourth part (S) having said first predetermined value, performing the steps of

(f.1) selecting a new object address (AB) in said second virtual address space such that $B=R$, and

(f.2) replacing said each internal object pointer (SPQR) with a new internal object pointer ($S_nP_nQ_nR_n$) such that $P_n=P$, $Q_n=A$, $R_n=R$ and said fourth part S_n equals a second predetermined value denoting completion of said relocation of said each object.

4. The method of claim 3 wherein said selecting step (f.1) comprises the steps of:

(f.1.1) if there exists an entry (P_iQ_i) in said object mapping table such that $P_i=P$, then selecting said new object address (AB) such that $A=Q_i$; otherwise

(f.1.2) adding a new entry (P_nQ_n) to said mapping table such that $P_n=P$ and $Q_n=A$.

5. The method of claim 3 or 4 wherein said performing step (f) is executed asynchronously with said replacing step (e).

6. In a multitasking network of one or more computer systems with distributed memory means for storing data and data structures, including a plurality of objects each created by some one of a plurality of concurrent processes each associated with at least one of a plurality of virtual address spaces in said distributed memory means, said each object being stored in an associated said virtual address space at an object address (AB) with a first part (A) and a second part (B), said each object including an object identifier (IJ) with a first part (I) and a second part (J), a data structure for denoting the global object identifier (IJ) and the local virtual object address (AB) of said each object, said data structure comprising:

a first part (P) representing said first part (I) of said global object identifier;

a second part (Q) representing said first part (A) of said local virtual object address; and

a third part (R) representing said second part (J) of said global object identifier, wherein said first and third parts (IJ) are held invariant over said network and $J=B$.

7. The data structure according to claim 6 to be used in a system for moving said each object to a second said process associated with a second said virtual address space, said system comprising:
 pointer generator means coupled to said associated virtual address space in said distributed memory means for combining said object identifier (IJ) and said base object address (A_0B_0) to produce a base object pointer ($P_0Q_0R_0$) with a first part ($P_0=I$), a second part ($Q_0=A_0$) and a third part ($R_0=J=B_0$):
 first linking means coupled to said pointer generator means in said distributed memory means for associating said base object pointer ($P_0Q_0R_0$) with said each object:
 address generator means coupled to said second process in said distributed memory means for selecting in said second virtual address space a first object address (A_1B_1) with a first part (A_1) and a second part (B_1) such that $B_1=B_0=J$:
 copying means coupled to said second process in said distributed memory means for copying said each object to produce a copied object at said second object address (A_1B_1) in said second virtual address space in said distributed memory means;
 and
 second linking means coupled to said copying means in said distributed memory means for associating a second object pointer ($P_1Q_1R_1$) with said copied object in said second virtual address space such that $P_1=P_0=I$, $Q_1=A_1$ and $R_1=R_0=J$.
8. The system of claim 7 further comprising:
 an object mapping table coupled to said each process in said associated virtual address space for storing an entry (P_iQ_i) representing the first part (P_i) and the second part (Q_i) for each said object pointer created by said each process:
 entry restoring means coupled to said object mapping table in said each associated virtual address space for selecting said second object address (A_1B_1) such that $A_1=Q_i$ if there exists an entry (P_iQ_i) in said object mapping table such that $P_i=P_0$; and
 entry creating means coupled to each object mapping table in said each associated virtual address space for adding a new entry (P_1Q_1) to said each mapping table such that $P_1=P_0$ and $Q_1=A_1$.
9. The system of claim 7 or 8 wherein said distributed memory means further includes a plurality of networked objects each including one or more internal object pointers (SPQR) referring to other said network objects, each said internal object pointer (SPQR) including a fourth part (S), said system further comprising:
 flagging means coupled to said each process in said distributed memory means for replacing said fourth part (S) of each said internal object pointer (SPQR) with a first predetermined value denoting pending

relocation of said each object to a new said process:
~~address selecting means coupled to said each process~~
 in said distributed memory means for selecting a new object address (AB) in said second virtual address space such that $B=R$ for each said internal object pointer (SPQR); and
 updating means coupled to said address selecting means for replacing said each internal object pointer (SPQR) with a new internal object pointer ($S_nP_nQ_nR_n$) such that $P_n=P$, $Q_n=A$, $R_n=R$ and said fourth part S_n equals a second predetermined value denoting completion of said relocation of said each object to said second virtual address space.




FIG. 1 (PRIOR ART)




FIG. 2 (PRIOR ART)




FIG. 5




FIG. 3A




FIG. 3B




FIG. 4A




FIG. 4B

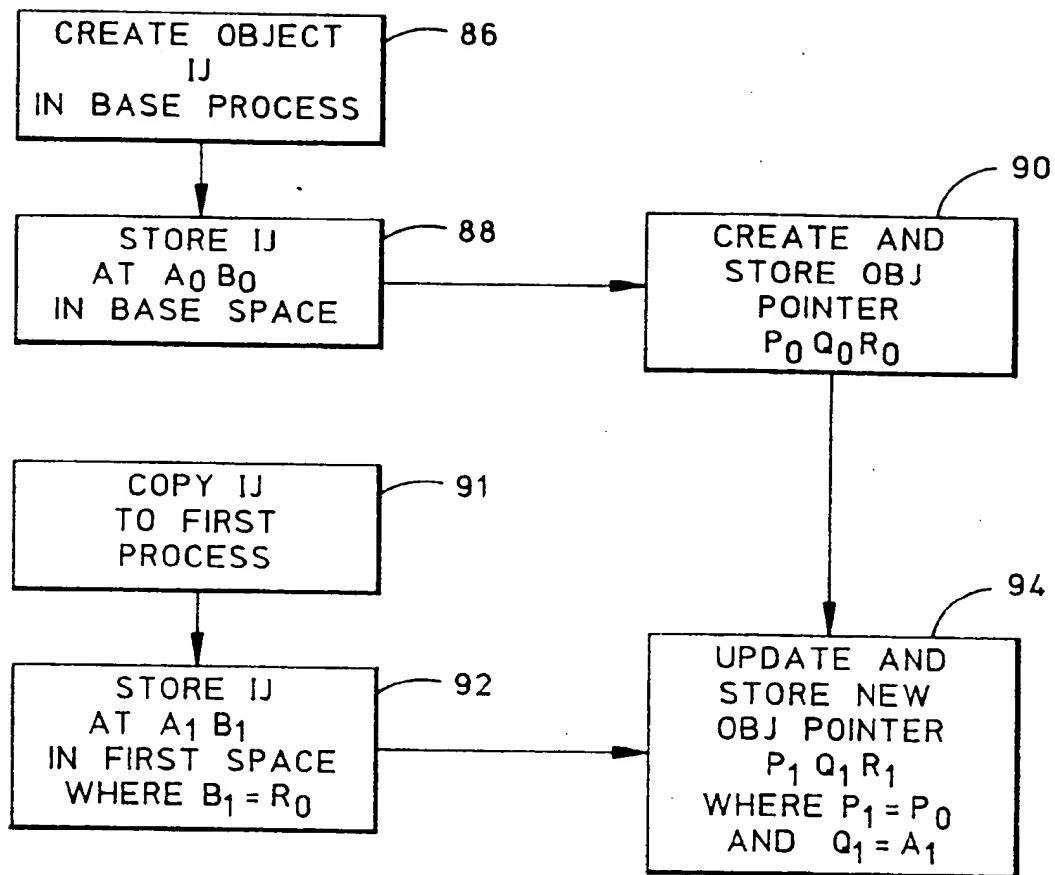


FIG. 6

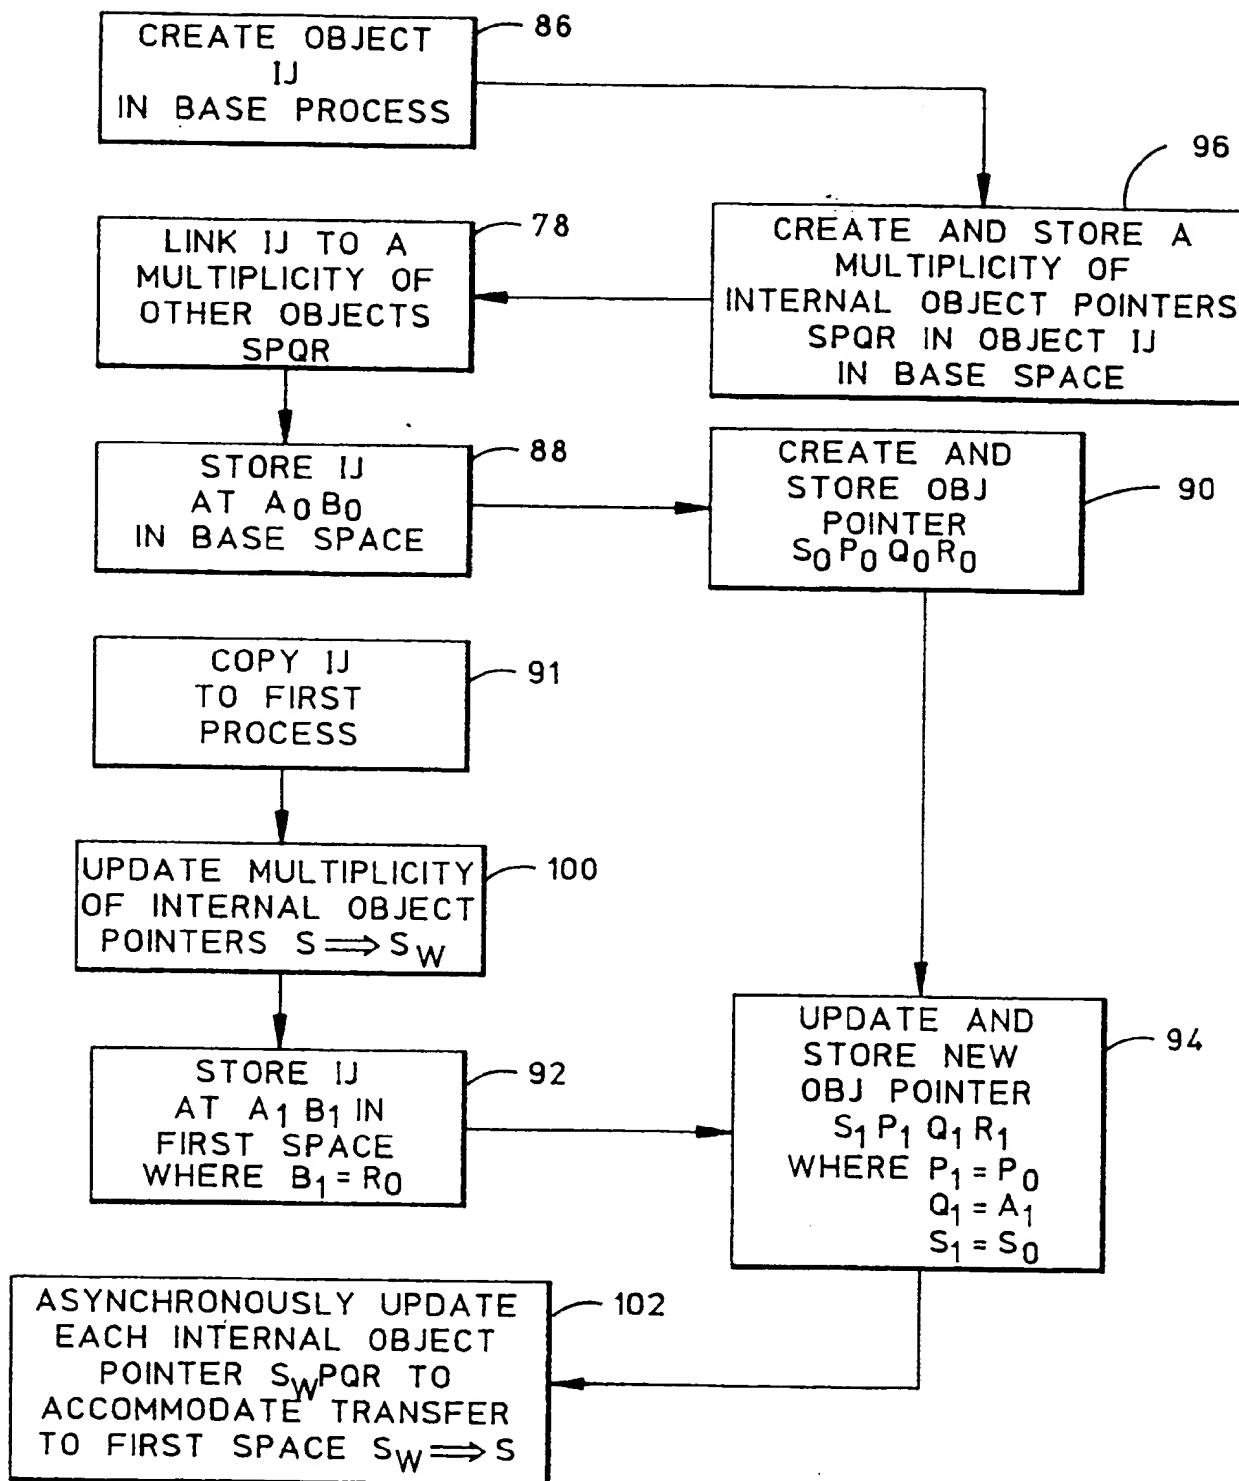


FIG. 7




FIG. 8

THIS PAGE BLANK (USPTO)

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 700 000 A3

(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
02.01.1997 Bulletin 1997/01

(51) Int Cl.⁶: **G06F 12/02**

(43) Date of publication A2:
06.03.1996 Bulletin 1996/10

(21) Application number: **95480098.3**(22) Date of filing: **24.07.1995**

(84) Designated Contracting States:
DE FR GB

(30) Priority: **30.08.1994 US 299042**

(71) Applicant: **International Business Machines Corporation**
Armonk, N.Y. 10504 (US)

(72) Inventors:
• **Munroe, Steven Jay**
Rochester, Minnesota 55901 (US)

• **Sirjani, Abolfazl**
Austin, Texas 78729 (US)
• **Voldal, Erik Edward**
Rochester, Minnesota 55901 (US)

(74) Representative: **Lattard, Nicole**
Compagnie IBM France
Département de Propriété Intellectuelle
06610 La Gaude (FR)

(54) **System and method combining a global object identifier with a local object address in a single object pointer**

(57) An object pointer data structure for efficiently combining an object identifier and an object address for use in object-oriented programming systems. An object address is a value that allows a client application or process to conduct high-performance operations on an object in the client's local virtual address space. An object identifier is a value that can be used to uniquely identify an object for the lifetime of that object across some defined domain, such as an entire universe of computer systems. The data structure of this invention defines an object pointer that is larger than the object address but smaller than the combination of the object identifier and object address. The truncated object pointer structure preserves all information from both object address and object identifier by forcing a portion of the local object address in each address space to be equal to a portion of the invariant object identifier. A local pointer mapping table may be used for efficiency in assigning local addresses to restored objects in each process.




FIG. 8



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 95 48 0098

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
A	US-A-4 660 142 (CLANCY GERALD F ET AL) 21 April 1987 * abstract *	1,6	G06F12/02
A	EP-A-0 542 483 (IBM) 19 May 1993 * abstract *	1,6	G06F
<p>TECHNICAL FIELDS SEARCHED (Int.Cl.6)</p> <p>G06F</p>			
<p>The present search report has been drawn up for all claims</p>			
Place of search		Date of completion of the search	Examiner
THE HAGUE		25 October 1996	Ledrut, P
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>			